

Reed, David P.; Schnitzer, Jason; Sekar, Dhanavikram

Conference Paper

Benchmarking Internet Progress: Analysis of FCC's Broadband Performance Data

ITS 33rd European Conference 2025: "Digital innovation and transformation in uncertain times", Edinburgh, UK, 29th June – 1st July 2025

Provided in Cooperation with:

International Telecommunications Society (ITS)

Suggested Citation: Reed, David P.; Schnitzer, Jason; Sekar, Dhanavikram (2025) : Benchmarking Internet Progress: Analysis of FCC's Broadband Performance Data, ITS 33rd European Conference 2025: "Digital innovation and transformation in uncertain times", Edinburgh, UK, 29th June – 1st July 2025, International Telecommunications Society (ITS), Calgary

This Version is available at:

<https://hdl.handle.net/10419/331298>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Benchmarking Internet Progress:
Analysis of FCC's Broadband Performance Data

David P. Reed, Jason Schnitzer, and Dhanavikram Sekar
University of Colorado Boulder

Abstract

This paper presents a novel longitudinal analysis of U.S. broadband performance trends from 2011 to 2023 using the Measuring Broadband America (MBA) dataset created and managed by the Federal Communications Commission. By systematically analyzing over a decade of actual user measurements, we quantify the evolution of key performance indicators including download and upload speeds, latency under load, along with their interquartile ranges (IQRs), providing the most comprehensive, empirically grounded portrait of broadband quality change available in the US to date. For example, we document the cumulative annual growth rates (CAGR) of median broadband service download speed of 20% (1.6 Mbps) for DSL, 30% (22 Mbps) for Cable, and 28% (28 Mbps) for Fiber systems between 2011 - 2023. Notably, between 2020 through July 2023, the CAGR in download speed for Fiber increased to 40% (73 Mbps) for Fiber. Our analysis indicates there may be some increase IQR for downstream and upstream speed across all DSL, Cable, and Fiber technologies over the 13-year study, meaning that there may have been an increase in the variability of broadband speeds across the middle 50% of users in the MBA program. We discuss the policy implications of these and other findings and make a concluding argument that as the FCC's MBA program was sunset in 2023, there is a strong need for the creation of a new open data platform to ensure that future broadband policy remains transparent, data-driven, and accountable. This study's data products and software artifacts are available to the research community at <https://github.com/UCBoulder/bclear/tree/main>.

Keywords: Broadband performance measurement, Internet performance measurement, broadband speed, broadband latency, longitudinal data analysis, Digital Divide

Table of Contents

<i>Abstract.....</i>	<i>2</i>
<i>I. INTRODUCTION.....</i>	<i>4</i>
<i>II. DATA DESCRIPTION AND VALIDATION</i>	<i>5</i>
Broadband Performance Tests	6
Data Validation.....	7
Enrichment	9
<i>III. EMPIRICAL ANALYSIS OF FCC'S MBA DATA.....</i>	<i>10</i>
Idle Latency Versus Download/Upload Speed.....	11
Downstream/Upstream Latency Under Load Versus Download/Upload Speed.....	12
Downstream Latency Under Load and Speed vs. Time	15
Upstream Latency under Load and Speed vs. Time	20
Low Performance Quartiles for Speed, Idle Latency, and LUL.....	21
Other Performance Metrics of Interest (Jitter).....	27
<i>IV. TECHNOLOGY POLICY IMPLICATIONS</i>	<i>29</i>
Insights from Longitudinal Analysis	29
Insights Into the Relationship between Speed and Latency	32
Broadband Policy Implications.....	34

Benchmarking Internet Progress:
Analysis of FCC's Broadband Performance Data

I. INTRODUCTION

As policymakers work to expand broadband infrastructure and close the Digital Divide, governments increasingly rely on broadband performance measurement as a cornerstone of data-driven policy. In the United States (US), the Federal Communications Commission's (FCC) *Measuring Broadband America* (MBA) program systematically collected broadband performance data from 2011 to 2023, capturing key metrics such as service speed and latency.¹ Despite its importance, the program was paused in July 2023 with no announced plans to continue collecting similar performance data in the future.

This paper presents a comprehensive longitudinal analysis of broadband service speed and latency trends in the U.S. using the complete MBA dataset. We assess how these trends align with evolving performance benchmarks and explore the implications of maintaining open access to such data for both researchers and policymakers. These questions are particularly timely as governments—including the US—increasingly tie broadband subsidies and universal service obligations to measurable performance standards. However, these standards often incorporate inconsistent or competing definitions of broadband service, with varying thresholds for speed, latency, and data allowances—all of which shape users' quality of experience (QoE).²

¹ This data collection effort implemented an active measurement methodology based on the SamKnows broadband performance data platform. See [Measuring Broadband America page \(accessed on June 9, 2025\)](#) for the official government archive of the 13 MBA Reports that were released by the FCC as well as to access the raw and validated broadband performance data collected by the program. For more on the Samknows platform (as well as Cisco Systems, its successor in interest), see <https://samknows.com/>.

² See D. Reed, *Is Speed Enough? Examining the Definition of Broadband and Its Implications for Public Policy* (August 3, 2021). TPRC49: The 49th Research Conference on Communication, Information and Internet Policy, Accessed on 09/06/2025 at <http://dx.doi.org/10.2139/ssrn.3898327>.

Beyond speed and latency, the MBA dataset contains additional underexamined metrics—including Latency under Load, jitter, and packet loss—that also affect QoE but remain largely absent from FCC reporting and underexplored in the research literature. While our primary focus is on service speed and network latency, we provide brief coverage of jitter performance as well. This study extends our earlier work by applying a reproducible data processing pipeline to the entire 138-month MBA dataset, offering new insights into long-term performance trends across these critical metrics.

We further compare the performance of three major U.S. wireline broadband technologies—DSL, cable, and fiber—against emerging definitions and standards for speed and latency.³ This analysis highlights how different statistical approaches to summarizing performance can lead to significantly different interpretations of broadband adequacy and policy compliance. Specifically, we evaluate how evolving benchmark definitions align—or misalign—with actual MBA-measured outcomes.

The remainder of the paper is organized as follows: Section II describes the MBA dataset and outlines our data validation and analysis methodology. Section III presents the results of our longitudinal performance analysis, and Section IV discusses the broader technology policy implications of these findings.

II. DATA DESCRIPTION AND VALIDATION

This section describes the foundation for the empirical analysis conducted in this study by describing how the FCC’s *Measuring Broadband America* (MBA) data was collected, filtered, validated, and enriched. The MBA dataset includes over 3.5 billion broadband performance

³ The MBA collects and reports results by the Internet Service Providers (ISPs) who participate in the program. Our research question is not pertinent to the trends of any particular ISP but on the general performance trends of the collective of all ISPs participating in the MBA program organized by the type of access network technology employed to provide broadband services.

measurements collected from thousands of Whitebox units deployed in U.S. households between February 2011 to July 2023.⁴ These measurements cover idle latency, latency under load, and download/upload speed. The data was processed through a structured pipeline consisting of three stages: unit identification, data filtering, and enrichment.

Readers interested in further discussion of our development of a data pipeline to ingest and analyze MBA program data are directed to an earlier paper that provides more technical detail.⁵

Broadband Performance Tests

Each Whitebox performed standardized tests according to FCC-prescribed methodologies. In this paper we analyze the following standardized broadband performance tests performed within the MBA program:

- **Download and Upload Speed.** Download and upload speed tests used HTTP-based throughput measurements across multiple concurrent TCP connections to emulate realistic user experiences. Download tests measure the flow of packets from remote test servers in the network to the Whitebox, while upload tests assess data transfer in the reverse direction from the Whitebox to the test server.
- **Idle Latency.** Idle latency tests are conducted when there is an absence of any other network traffic present at the Whitebox location in the home. This “quiescent”

⁴ After a call for volunteers for testing from the MBA program, SamKnows selected 12,000 participants to receive a “Whitebox” or off-the-shelf router with tests embedded in the software. For a detailed overview of the SamKnows platform and open methodology employed by the FCC, see FCC Website Measuring Broadband America – Open Methodology, accessed on June 15, 2025 at <https://www.fcc.gov/oet/mba/measuring-broadband-america-open-methodology>.

⁵ Schnitzer, Reed and Gandhi, *op cit.*, Section III. We have made some minor technical changes to the data pipeline since this work was published in 2023. These include: 1) implementing performance improvements through structural changes (vectorized calculations instead of traditional for loops), 2) moved to a different Python dataframe – Polars instead of Pandas – in the backend for 10x faster data processing, and 3) moved to Apache Parquet data – a more efficient data storage format – for almost 75% reduction in data storage requirements.

network state provides a measure of the baseline delay of the network. This test used 64-byte UDP packets sent once per second for a 20-second test window. Note that idle does not refer to the state of the broadband access network traffic load during the time of test execution, but only to the subscriber network in which the Whitebox test agent resides.

- **Latency Under Load (LUL).** LUL tests are conducted in the presence of some other significant amount of network activity, and is, therefore, considered to give a more representative of real-world performance or “working” latency that captures network behavior during congestion. The LUL test sends UDP packets every 100 milliseconds concurrently with upload and download speed tests, thus simulating real-time application conditions such as video calls or gaming.

Data Validation

This section provides a high-level overview of our methodology to validate the FCC MBA latency measurement data. Validation is necessary since the raw data collected in the measurement effort can contain erroneous test results or otherwise unverified information. Our study uses the data collected from the broadband tests listed in the prior subsection.

Our data validation pipeline consists of three distinct stages.

1. **State 1: Resolve Measurement Topology and Access Technology.** In this first stage we establish the correlation between all `unit_id` participating in the MBA consumer panel and the Internet Service Provider (ISP), as well as the ISP networking technology (DSL, Cable, Fiber). These `unit_id` are randomly generated, which preserves the volunteer panelists’ anonymity.

2. **State 2: Validation Filters.** The next stage applies filters to the data that are designed to align with the validation criteria utilized by the FCC. The filters applied include excluding test events with zero entries for all columns, excluding speed test events where speeds were recorded above 940 Mbps,⁶ excluding test events where round trip times exceeded 3000 ms or were less than 0.05 ms, and excluded test instances with less than 50 successful packets or where packet loss exceeded 10%.
3. **State 3: Validated Data Amendment.** After the data filtering process, we enhance the raw dataset by amending four new fields:
 - a. *test type*: Denotes whether a test is categorized as an “on-net” or “off-net” test using the procedure described above. For this purpose, on-net refers to the test occurring on a test node located on a specific ISP’s network and off-net refers to the test occurring on a test node located outside a specific ISP’s network. Test nodes located on a particular ISP’s network (“on-net” test nodes) might introduce bias with respect to the ISP’s own network performance.⁷
 - b. *validation type*: Indicates whether the `unit_id` in the raw data is part of the FCC’s validated units list. If the `unit_id` is found in the *unit profile.csv*

⁶ The FCC results for Fiber had to “normalize” the fiber speed measurements for 1 Gbps service to 940 Mbps since this was the maximum speed at the transport layer of the Samknows Whitebox v8.0 used in the speed tests. We applied this exclusion filter to any test events above 940 Mbps in the upstream or downstream as being erroneous for this reason.

⁷ See FCC’s Office of Engineering Technology, Measuring Broadband America: Technical Appendix to the Thirteenth MBA Report, p. 62. Accessed on June 16, 2025 at <https://data.fcc.gov/download/measuring-broadband-america/2023/Technical-Appendix-fixed-2023.pdf>.

files, it is classified as “fcc”. Otherwise, it is labeled as “bmc” (representing excluded and “off-net” Whiteboxes).

- c. *dtime local*: Local time conversion using *timezone_offset* and *timezone_offset_dtc* fields from the *unit_profile.csv* and *excluded_units.csv* files. If the time falls within the daylight-saving period, it is adjusted to the local daylight-saving time zone. This field plays a crucial role in categorizing the “time category” field, described below.
- d. *time category*: This field classifies timestamps into one of three parts of the day used by the FCC syntax: “peak hours” (7PM-11PM weekdays), “off-peak hours” (weekdays), or “satsun” (all weekend hours).

Enrichment

Upon the application of all filters, these fields are merged with the *unit_id_map.csv* file to create eight new fields to the original dataset: 1) resolved ISP name, 2) ISP access network technology type, 3) advertised download speed tier, 4) advertised upload speed tier, 5) test type (on-net/off-net), 6) inclusion on FCC validated unit list, 7) timestamp adjusted to local time zone, and 8) time-of-day classification into Weekday Peak, Weekday Off-Peak, or Weekend.

To align our analysis with the FCC’s analytical methodology, we take one final step before the validated data products of upload/download speed, idle latency, and LUL performance metrics are ready for longitudinal analysis. This step takes the remaining validated data after filtering and computes the average value of the performance metric for each Whitebox on a 30-day basis. Establishes equal treatment of each client to avoid bias favoring more frequently reporting Whiteboxes (e.g., median speed results for fiber technology reflect the median of all

average values of Whiteboxes connected to fiber networks)⁸ This establishes equal standing of each Whitebox or unit_id to avoid bias that could occur by having the values generating from more frequently reporting Whiteboxes contributing more to statistical calculations of median or percentile determinations (*e.g.*, median LUL results for fiber technology reflect the median of all average 30-day values of Whiteboxes connected to fiber networks).

In summary, the information provided in this paper section provides guidance for implementing a reproducible framework for transforming the raw FCC MBA data into an analyzable dataset. The methodology for developing this data pipeline is grounded in FCC validation procedures as reported by the FCC. The establishment of this data pipeline is to enable fine-grained temporal, technological, and provider-level analysis of broadband performance metrics included in the MBA program. As will be demonstrated in the remainder of the paper, the platform supports the longitudinal comparisons necessary for tracking broadband trends and technological policy implications. In short, by combining the measurement rigor of the MBA program with data enrichment provided by the data pipeline to access powerful data visualization software, the data pipeline creates an opportunity to access one of the most comprehensive broadband performance datasets available for research, policy evaluation, and public accountability.

III. EMPIRICAL ANALYSIS OF FCC'S MBA DATA

This section reports our analysis of the MBA program data to identify trends in the performance of broadband services over time. Using validated MBA data as described in Section II, we now review how key broadband performance metrics have trended spanning 2011 through

⁸ This step was not performed in our prior work cited earlier and therefore could result in small differences in our reported results for idle latency and LUL results. See Schnitzer, Reed, and Gandhi, *op cit*, Section IV.

2023. The performance metrics that we examine in this section are measured download/upload speed, idle latency, latency under load, and jitter. We begin our analysis by looking at speed and latency in integrated fashion by plotting the changes in these two metrics over time on a combined graph to better understand and illustrate any dependencies that may exist between these two important attributes of broadband service.⁹

To begin this review, we first look at idle latency versus download/upload speed and downstream/upstream LUL versus download/upload speed. We then report our analysis of download/upload speed versus LUL. The section then closes with a brief longitudinal analysis of the jitter performance data.

Idle Latency Versus Download/Upload Speed

The boxplot in Figure III.1 plots the idle latency versus download/upload broadband speed for 2023, which in the case of MBA measurements is represented by data collected between January thru July 2023 for all technology types. For this graph, the 30-day measured speed value determines which speed category the Whitebox falls into, with the bins representing increasing categories of speed ranges from left to right. The size of the box is set by the interquartile range (IQR) representing by the 75th quartile (top of the box), 25th quartile (bottom of the box), and 50th percentile (median black line inside the box) of the 30-day idle latency measurements obtained for the same Whitebox. The vertical bar in each bin shows the maximum and minimum values of idle latency for all measured speed events falling into this bin. The IQR provides a measure of the dispersion or the magnitude of the spread in latency measurements within each bin. As idle latency represents the underlying time for transport in the presence of

⁹ Due to significant gaps over time in the MBA dataset of entries for the advertised speeds of the service tier of the household, our analysis only considers measured service speeds which we refer to simply as “speed” in the analysis.

little or no network congestion (a quiescent state), it theoretically should establish the best case (lowest expected) latency for the network connection established between the test server and the client device (the Samknows Whitebox in the household).¹⁰

The results clearly demonstrate an expected relationship between latency and speed. These plots demonstrate a decrease in median idle latency from roughly 35 ms to 15 ms when *downstream* speeds increase from below 50 Mbps to above 500 Mbps.¹¹ There also is a decrease in median idle latency from about 25 ms to 12 ms when *upstream* speeds increase from below 50 Mbps to above 500 Mbps. Thus, the results show that as broadband download/upload speeds increase from <50 Mbps to 500+ Mbps, median idle latency decreases over 50% since the time to transmit and process each Internet packet is reduced as service speed increases. Because network engineers and policy makers are usually most interested in network performance in the presence of some amount of congestion, we will not spend further effort here to analyze these results any further.

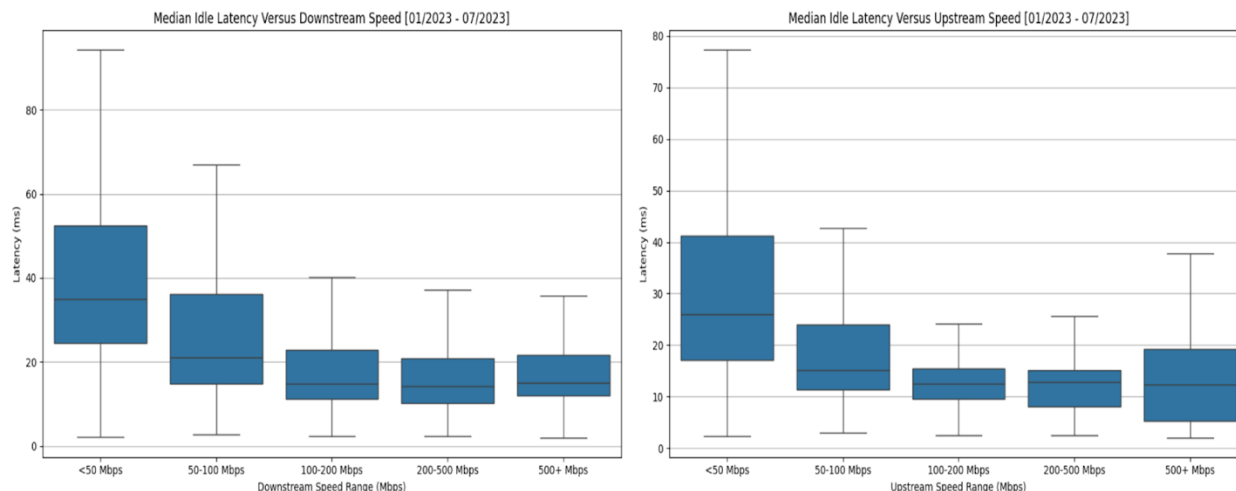
Downstream/Upstream Latency Under Load Versus Download/Upload Speed

The boxplot in Figure III.2 plots of downstream/upstream LUL versus download/upload broadband speed for FY2023. This graph uses the same speed bins as Figure III.1 and includes all technologies combined into each boxplot.

¹⁰ The three basic contributors to latency are packet queuing due to congestion, packet processing in network routing equipment, and transmission speed. With idle latency, congestion is removed or minimized leaving packet processing and service speed as the major components remaining to set the baseline level.

¹¹ The FCC's idle latency test measures the round-trip time between the test server and Whitebox when the subscriber network is idle and therefore is not a separate measure of latency in the downstream or upstream direction. The downstream and upstream data shown in Figure III.1 is associated with the direction of the speed data.

Figure III.1: Idle Latency (Quiescent) vs. Downstream/Upstream Speed FY2023

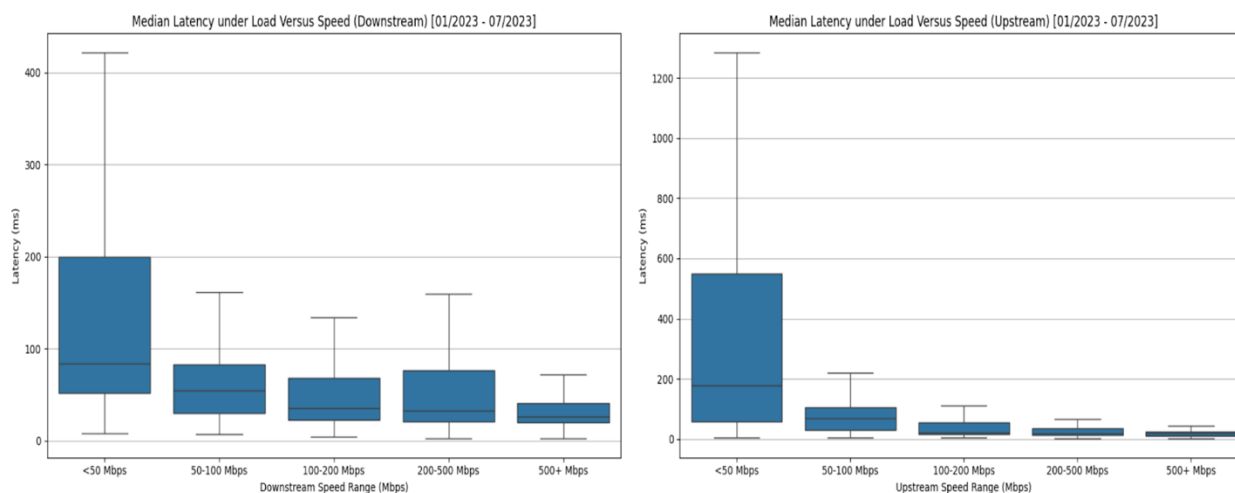


This is a more interesting case since it shows network performance under stress by injecting congestion into the test methodology for latency. Now the plots demonstrate a decrease in median *downstream* LuL from roughly 90 ms to 25 ms when speeds increase from below 50 Mbps to above 500 Mbps. The decrease in median *upstream* LuL is from about 190 ms to 25 ms when speeds increase from below 50 Mbps to above 500 Mbps. These results show that as broadband download/upload speeds increase from <50 Mbps to 500+ Mbps, median LUL decreases 70% in the downstream and 85% in the upstream.

These results are in line with expectations. First, the inverse relationship between speed and latency is again evident as the median values for LuL fall off in exponential fashion as the broadband speeds increase from below 50 Mbps to more than 500 Mbps. Second, the expected impact of congestion created by the LuL test metric itself is illustrated as well. The median LuL values between idle latency (Figure III.1) and LuL (Figure III.2) reflect a rough increase of 40% to 60% in downstream latency and 15% in upstream latency over the range of service speeds between 50 Mbps and over 500 Mbps. This demonstrates the impact of network congestion

within the Samknows platform used by the FCC for the MBA program.¹² Third, the boxplot format clearly shows more dispersion in the latency results in lower broadband speed categories than the higher speed ones. This might seem counterintuitive to expectations but there is a straightforward explanation derived from the recognition that these plots include all the access network technologies employed in FY2023. Most of the data obtained for the highest speed categories is obtained from advanced fiber or recently upgraded cable networks, while the lowest speed category reflects a considerable amount of DSL technology, which has from much higher latency figures that can vary in unpredictable fashion based upon the length of copper twisted pairs.

Figure III.2: Latency under Load (Working Latency) vs. Downstream/Upstream Speed FY2023



¹² Stated differently, this impact of congestion on latency cannot be generalized to be representative to overall Internet performance on the whole. The Samknows platform measurements are accurate for the specific network topology that existed between the Samknows test servers and the client Whiteboxes. This topology includes portions of the Broadband Internet Access Service provider network necessary to access the Samknows server (e.g., local access network, metropolitan area network or middle-mile network, perhaps even portions of a backbone network). In addition, the topology over which Samknows collected broadband performance data is known to have changed over this period of time as well, which would have changed the measured impact of congestion on the LuL tests.

Figure III.3 and Figure III.4 provide full 13-year longitudinal plots of LUL and speed versus time on the same graph for the technology types of DSL, Cable, and Fiber. Each chart represents the following information:

- For each technology type, a graph with an x-axis that depicts time as the independent variable between 2011 – 2023 and a y-axis that shows the dependent variables of LUL (in ms on the left-hand side) and service speed (in Mbps on the right-hand side).
- Plots of the median line and IQR (shaded blue) for the measured values of LUL over time. Median LUL and other percentiles are computed within 30-day windows spanning the entire 13 years by technology type.
- Plots of the median line and IQR (shaded orange) for the measured values of speed over time. Median speed and other percentiles are computed within 30-day windows spanning the entire 13 years by technology type.

The discussion that follows for these charts is divided into downstream and upstream sub-sections, beginning with the downstream data.

Downstream Latency Under Load and Speed vs. Time

Figure III.3 shows the graph of downstream LuL and speed versus time by technology type for the local access networks based on Fiber, Cable, and DSL technologies. Fiber stands out with the highest median download speed ranging from 25 Mbps in 2011 to 375 Mbps in 2023, while DSL has the lowest speed in the range of 3 to 23 Mbps during this time, and Cable falls in between with a range of 15 to 300 Mbps during this time.¹³ Although not depicted in Figure III.1,

¹³ This increase in speed cannot be interpreted or projected as being representative of the overall growth in speed for the Internet for reasons that will be discussed further in the Section IV. It does provide guidance on the individual performance improvements experienced by users on specific technologies. The distribution of households contributing data to the MBA program is not representative of the actual adoption of each service speed tier in the

an exponential trend line derived from these findings suggest an annual improvement in the median download speed of approximately 1.6 Mbps for DSL, 22 Mbps for Cable, and 28 Mbps for Fiber systems between 2011 - 2023. Since 2020 through July 2023, the annual increase in download speed was substantially faster than in prior years 2011 – 2019, growing to annual increases of 3.2 Mbps for DSL, 43 Mbps for Cable, and 75 Mbps for Fiber.

The median downstream LuL has decreased between 2011 – 2023 from 128 to 66 ms for DSL, 83 to 41 ms for Cable, and 141 to 23 ms for Fiber. This reduction in downstream LuL over time translate into an average decrease of about 5, 3, and 9 ms per year between 2011 – 2023 for DSL, Cable, and Fiber systems, respectively. Since 2020 through July 2023, the annual decrease in median LuL accelerated faster than in prior years 2011 – 2019, representing annual decreases of 6 ms for DSL, 5 ms for Cable, and 5 ms for Fiber.¹⁴

The IQR results graphed in Figure III.3 for all of the technology types provide evidence that the dataset is right or positively skewed, meaning the difference between the 75th percentile and median is greater than the difference between the median and the 25th percentile. This indicates the tail to the 75th percentile is much longer than that of the 25th percentile. Consequently, the 75th percentile speed values will be higher (which would be considered a positive for the broadband performance) but the 75th percentile LUL values will be higher as well (which would be considered a negative for broadband performance).

This skew contributes to a large spread and variability in this data between the 25th and 75th percentiles in downstream speed. Figure III.3 shows that the IQR has increased over time for

broadband Internet access market, nor the market share of this market by each technology type. To make it representative, a weighting for each contributor needs to be applied reflecting the percentage of households subscribing to the speed tier and technology type of the contributing household. We do not attempt to apply such a weighting in our analysis of the results.

¹⁴ These estimates assume the average decrease in the dispersion for IQR LuL per year between 2011 – 2019 was 3 ms for DSL, 2 ms for Cable, and 10 ms for Fiber.

speed measurements and decreased over time for LuL measurements between 2011 – 2023. The actual gap between the 25th and 75th percentiles in downstream speed increased between 2011 to 2023 from 3.9 to 45 Mbps for DSL, 6 to 338 Mbps for Cable, and 9 to 438 Mbps for Fiber, respectively. As will be discussed in Section V, there are good reasons to question the magnitude of this rise in the IQR spread over time.

Figure III.3 also illustrates that the spread between the 25th and 75th percentiles in downstream LuL narrowed between 2011 – 2023 by between 60% - 65% for all three technologies.¹⁵ This decrease in LUL across all three technologies reflects a steady improvement during this time in the consistency and reliability of broadband performance during periods of network stress.

¹⁵ This estimate assumes the spread in the downstream IQR for LuL between 2011 and 2023, respectively, was 241 ms and 92 ms for DSL, 164 ms and 63 ms for Cable, and 105 ms and 36 ms for Fiber.

Figure III.3: Longitudinal Downstream Latency under Load vs. Speed (DSL, Cable, Fiber)

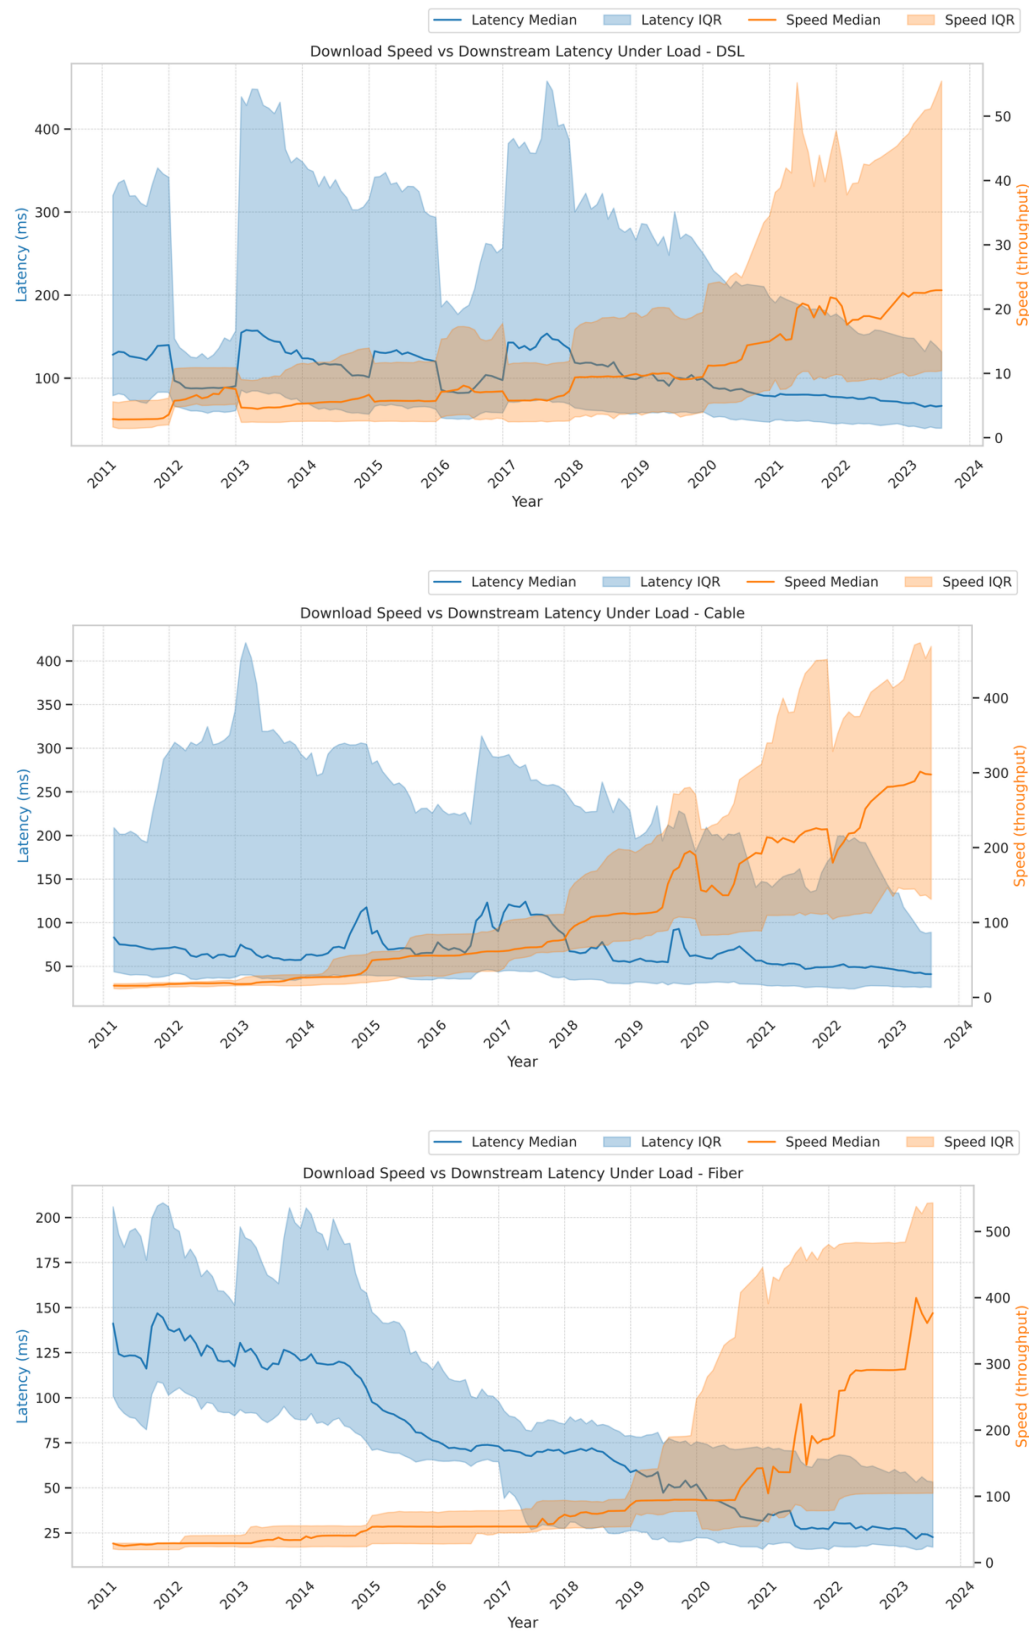
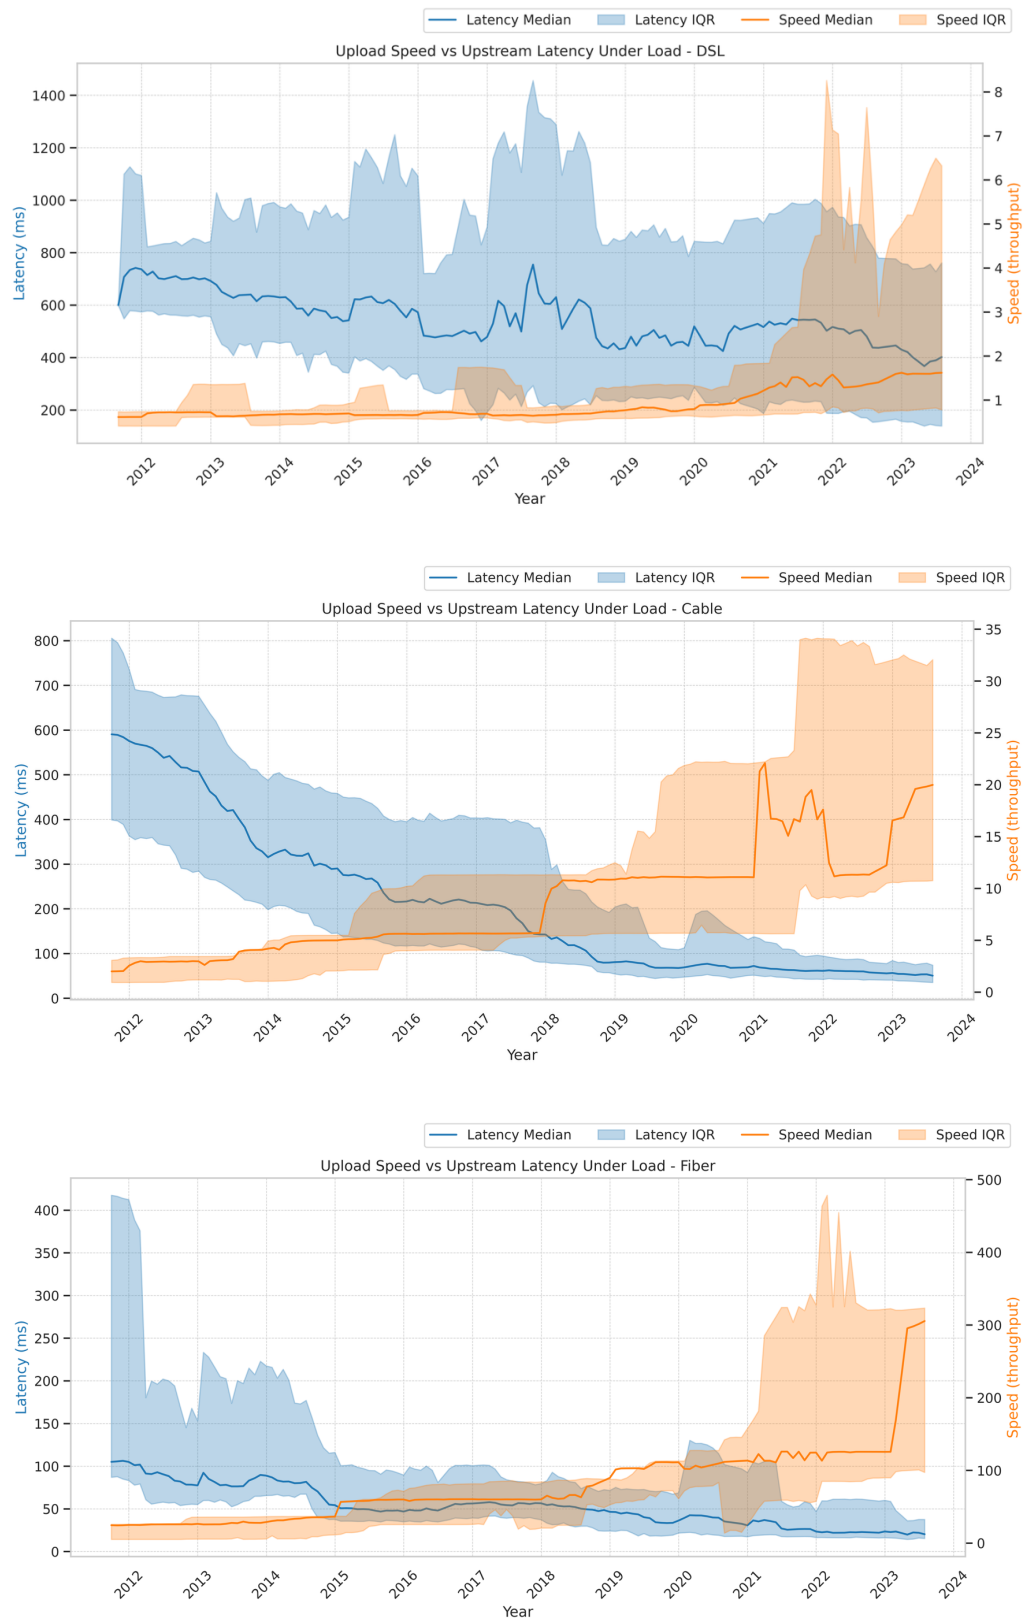


Figure III.4: Longitudinal Upstream Latency under Load vs. Speed (DSL, Cable, Fiber)



Upstream Latency under Load and Speed vs. Time

Figure III.4 shows the graph of upstream LuL and speed versus time by technology type. Fiber again exhibits the highest median upload speeds, ranging between 24 Mbps in 2011 to 305 Mbps in 2023, surpassing both DSL's range of 0.6 to 1.6 Mbps and Cable's range of 1.8 to 20 Mbps. This translates into an increase of 0.1, 1.5, and 30 Mbps per year for DSL, Cable, and Fiber systems, respectively. Noteworthy observation for this data is that Fiber demonstrated substantial upstream speed enhancements of roughly 50 Mbps in 2018 and 200 Mbps in 2023. Cable had two more modest increases in upstream speed of 7 Mbps in 2018 and 10 Mbps in 2021 (though this latter boost was uneven through to 2023). DSL upstream speeds remained relatively constant values slowly increasing from 1 to 2 Mbps during the full duration of the 13-year study.

The median upstream LuL has decreased between 2011 – 2023 for this set of households from 707 to 400 ms for DSL, 590 to 50 ms for Cable, and 105 to 20 ms for Fiber. This improvement in upstream LuL translates into an average rate of about 24, 43, and 7 ms per year between 2011 – 2023 for DSL, Cable, and Fiber systems, respectively. The Cable and DSL figures are much larger than downstream LuL improvement due to high initial starting points of 590 and 707 ms for the LuL measurements for these technologies.

The IQR results from Figure-III.4 show the dispersion between the 25th and 75th percentiles in upstream speed in 2023 were 5 Mbps for DSL, 23 Mbps for Cable, and 213 Mbps for Fiber. As noted earlier in the discussion of the IQR values for download speed, there are good reasons to believe that these IQR values are overstated, as will be discussed in Section V.

Figure III.4 also shows the upstream speed IQR increased between 2011 to 2023 by 1767% for DSL, about 995% for Cable, and about 914% for Fiber.¹⁶ It is interesting that the

¹⁶ This estimate assumes the spread in upstream IQR for 2011 and 2023, respectively, was 0.3 Mbps and 5.6 Mbps for DSL, 2.1 Mbps and 23 Mbps for Cable, and 21 Mbps and 213 Mbps for Fiber.

increase in upstream speed IQR is much lower (by a factor of about 5) for the upstream as compared to the downstream . For cable this outcome can be explained by the high cost of adding upstream capacity versus downstream capacity on their network and kept the spread of bandwidth across the upstream speed tiers to a smaller amount. For Fiber, the explanation is less clear. Some of the ISPs deploying Fiber in the MBA program have tried to standardize around fewer service tiers, although further research on this topic is needed.

Contrary to expectations, the LuL IQR actually grew as well between 2011 – 2023 by about 13% for DSL. For Cable and Fiber, it narrowed or decreased during this time by 90% and 93%, respectively.¹⁷ This outcome contrasts substantially from the uniformity of change for downstream LuL across all technologies. The similarity between Cable and Fiber is to be expected given similar proportional increases in upstream speed identified in the discussion above that will serve to decrease latency measurements due to these higher speeds. For DSL, one hypothesis could be that the small magnitude of upstream speed changes over time preclude the possibility of significant decrease in any latency resulting due to service speed.

Low Performance Quartiles for Speed, Idle Latency, and LUL

Because there is much interest on the part of policy makers in the current lower percentiles of Internet performance, this section presents charts of the lowest performing quartiles for download and upload speed, idle latency, and downstream and upstream LUL. Each Figure provides the following four graphs using only the most current MBA data for the months of January – July 2023: 1) an All Technologies graph that aggregates DSL, Cable, and Fiber data, 2) a DSL graph, 3) a Cable graph, and 4) a Fiber graph. A brief explanation is provided for each Figure as well using the performance benchmarks of 25 Mbps and 100 Mbps for download speed

¹⁷ This estimate assumes the spread in upstream IQR for 2011 and 2023, respectively, was 552 ms and 622 ms for DSL, 400 ms and 39 ms for Cable, and 331 ms and 22 ms for Fiber.

(the prior and current speed thresholds used in the US), 20 Mbps for upload speed (current benchmark used in the US), and 95th percentile latency below 100 ms (a common benchmark in US broadband funding programs).

Figure III.5: Low Performance Download Speed Percentiles (Overall, DSL, Cable, Fiber)

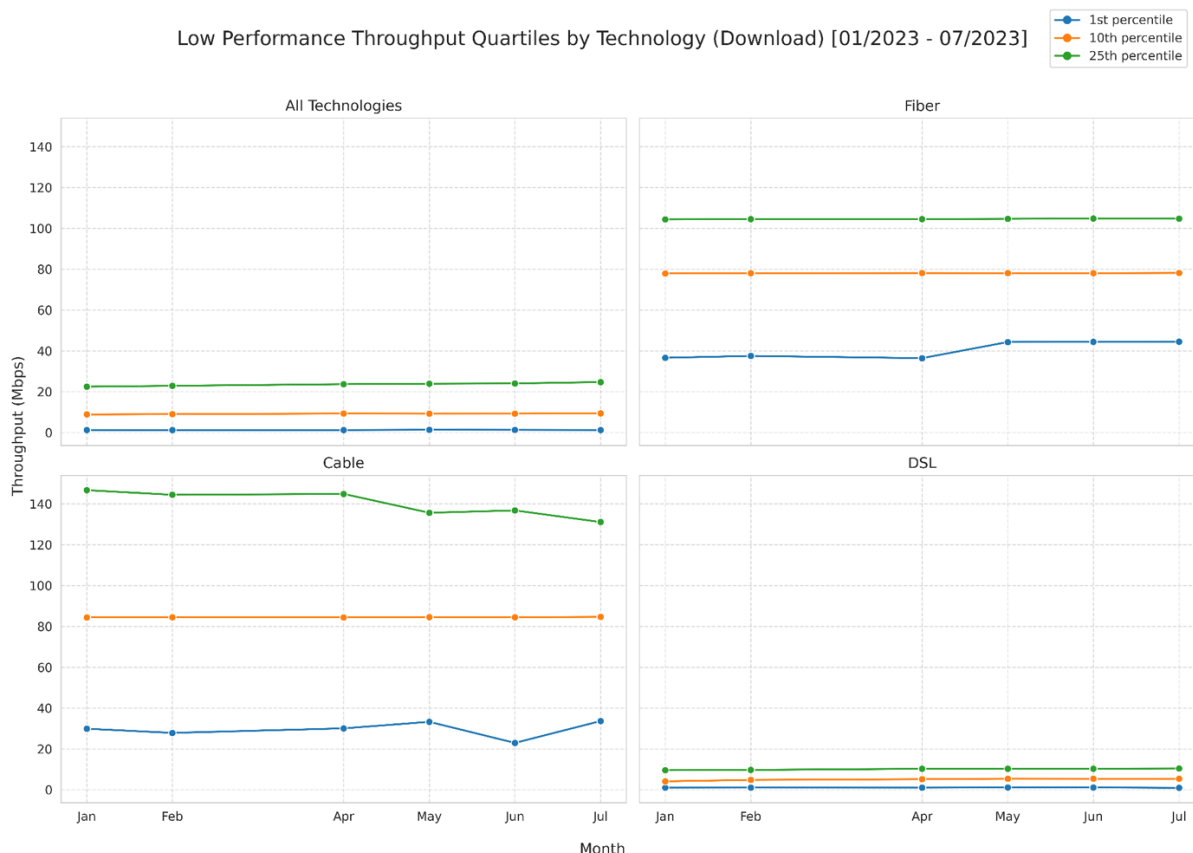


Figure III.5 shows the lower percentiles of download speed tests. These results indicate that DSL households falling in the 25th percentile on down do not meet a 25 Mbps threshold. All Cable and Fiber households falling in the 10th percentile on download speed meet a 25 Mbps threshold but the 25th percentile does not meet a 100 Mbps threshold. All Fiber households falling in the 10th percentile on download speed meet a 25 Mbps threshold but the 25th percentile does not meet a 100 Mbps threshold.

Figure III.6: Low Performance Upload Speed Percentiles (Overall, DSL, Cable, Fiber)

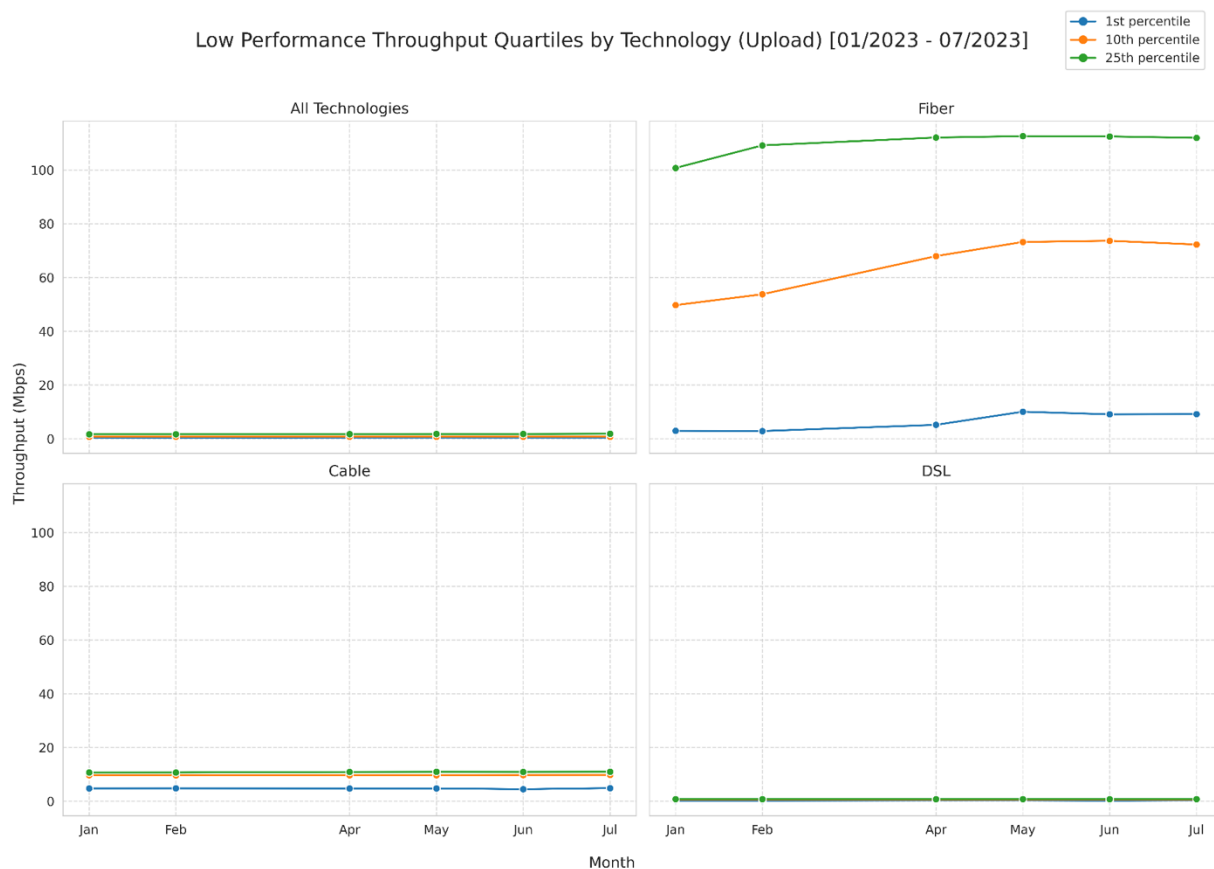


Figure III.6 shows the lower percentiles of upload speed tests. These results indicate that DSL and Cable households falling in the 25th percentile on download speed do not meet a 20 Mbps threshold. All Fiber households falling in the 1st percentile would not meet a 20 Mbps threshold while the 10th percentile on up would exceed the 20 Mbps threshold.

Figure III.7: Low Performance Idle Latency Percentiles (Overall, DSL, Cable, Fiber)

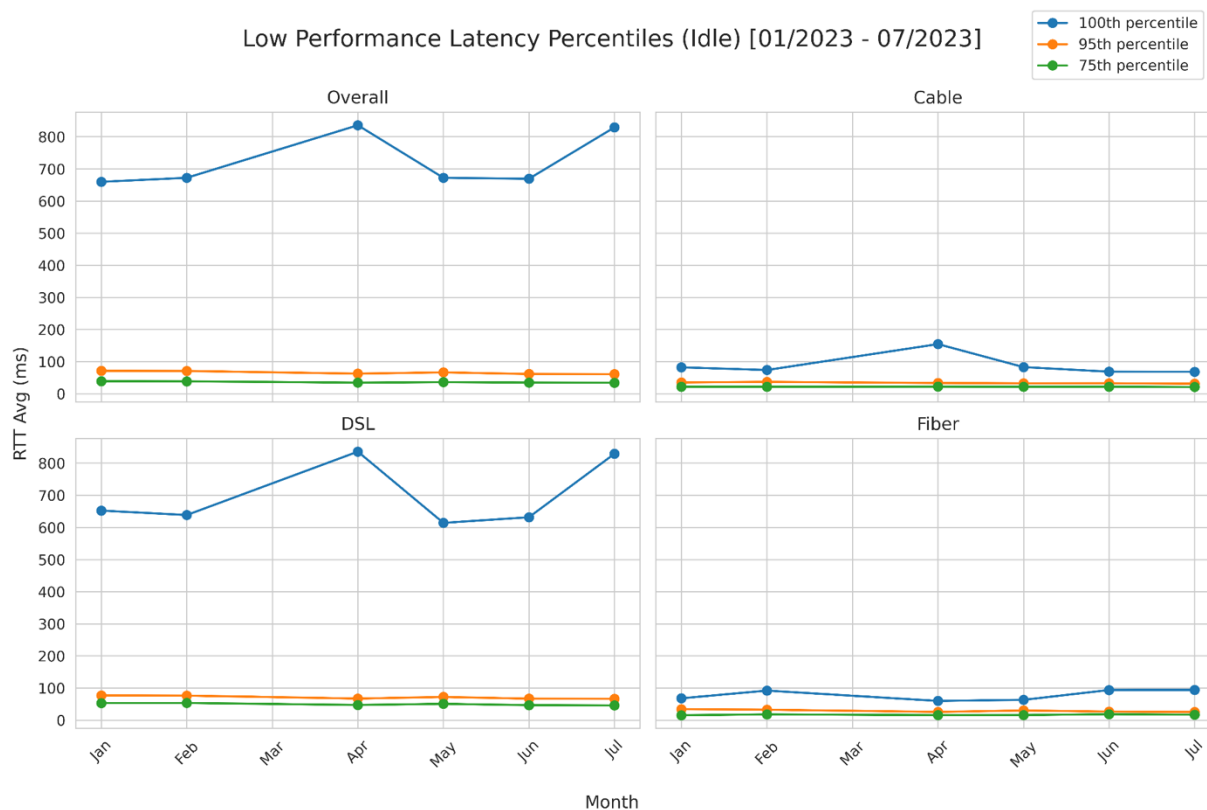


Figure III.7 shows the lower percentiles of idle latency tests. These results indicate that with the exception for the 99th percentile of DSL,¹⁸ and the 99th percentile of Cable in April, all other percentiles shown in the graph would meet a 100 ms threshold at the 95th percentile or below. Given idle latency represents the baseline latency for the network without congestion present, the result that most low performance latency percentiles meet these thresholds indicates that the broadband infrastructure supporting the testing is designed to operate efficiently when congestion is minimized.

¹⁸ The FCC reports the results of idle latency as the 99th percentile, which accounts for the 100th percentile entry in the legend of this graph.

Figure III.8: Low Performance Downstream LUL Percentiles (Overall, DSL, Cable, Fiber)




Figure III.8 shows the lower percentiles of downstream LUL tests. These results indicate that Fiber meets the threshold of 95th percentile below 100 ms. Cable meets the 100 ms threshold partially with its 75th percentile results and DSL does not meet the 100 ms threshold at any percentile shown. These results demonstrate how the different network technologies react to downstream congestion. The 99th percentile results provide the extreme example of this reaction with the lowest 1% of test results. The resulting latency for this percentile is above 1100 ms for DSL, around 450 ms for Cable, and around 300 ms for Fiber. All of this amount of latency would result in unacceptable quality of real-time VoIP or video conferencing applications.

Figure III.9: Low Performance Upstream LUL Percentiles (Overall, DSL, Cable, Fiber)

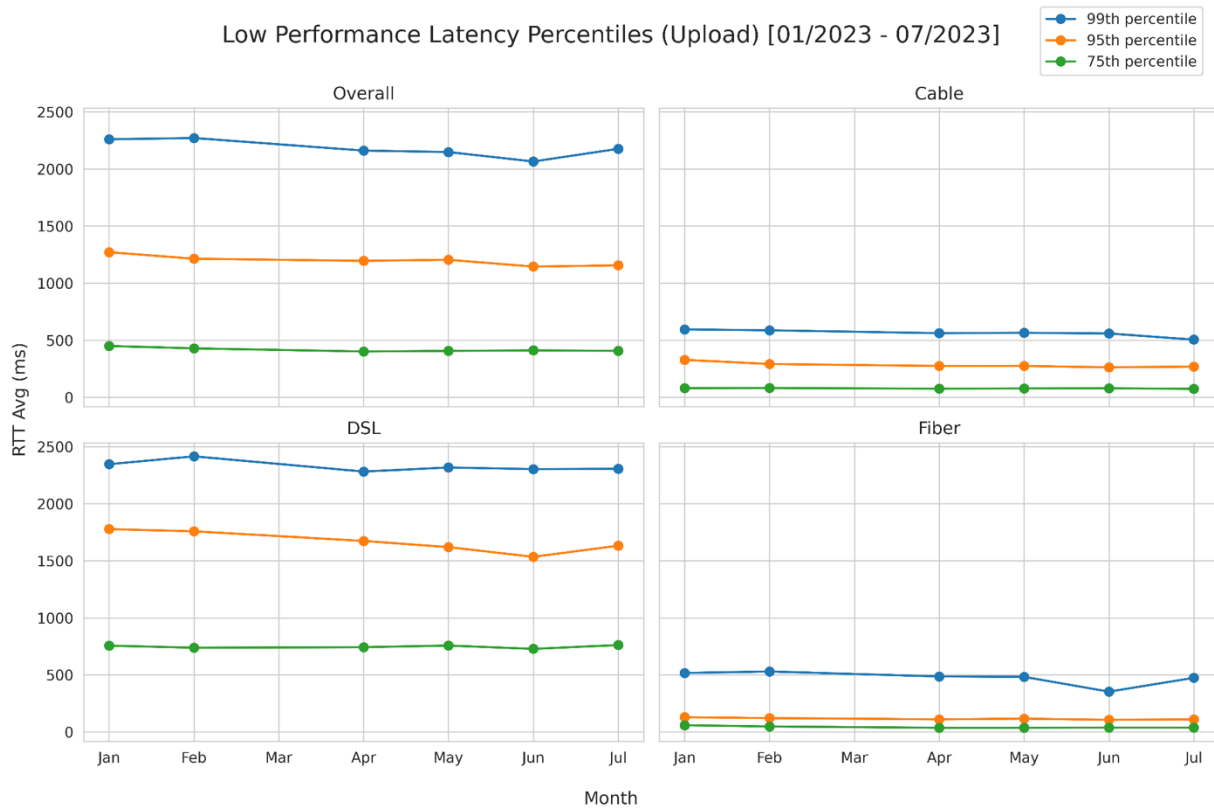


Figure III.9 shows the lower percentiles of upstream LUL tests. These results indicate that only Fiber partially meets a 100 ms threshold at the 95th percentile (and fully meets it at the 75th percentile). Cable meets the 100 ms threshold partially with its 75th percentile results and DSL does not meet the 100 ms threshold at any percentile shown. These results show how the different network technologies react to upstream congestion, which can be more of a reaction than downstream since the upstream capacity is often smaller than available in the downstream. Again the 99th percentile results provide the extreme example of this reaction with the lowest 1% of test results. The resulting upstream latency for this percentile is above 2400 ms for DSL, above 500 ms for Cable, and just under 500 ms for Fiber. As with the downstream, this amount of latency would result in unacceptable quality of real-time VoIP or video conferencing applications.

Other Performance Metrics of Interest (Jitter)

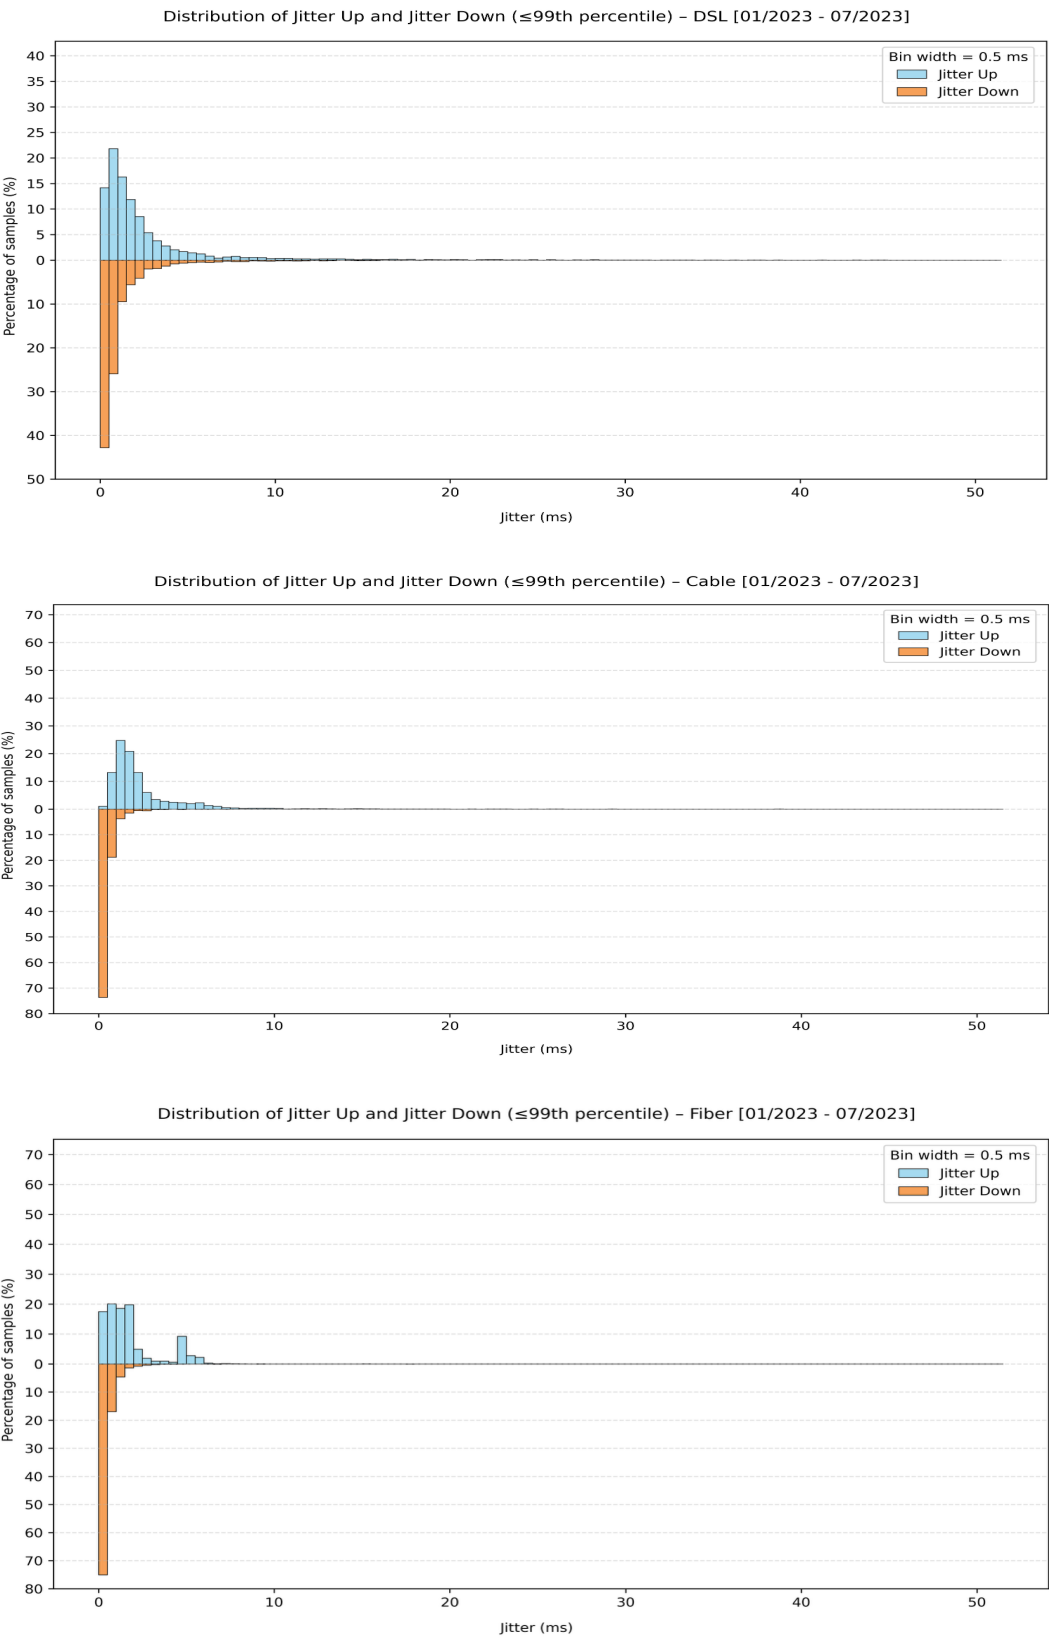
The MBA program collects data on several other performance measures including jitter which can provide a strong characterization or implications for the overall quality of real-time applications. The FCC defines jitter as the variation in the time between data packets arriving to the client. High variation in this timing between packets is often caused by network congestion.¹⁹ Jitter is an important performance metric for real-time applications such as voice over IP (VoIP) or video conferencing, where timing consistency is critical. High jitter can result in a poor quality of experience (QoE) for the user even if there is adequate capacity in the network.

We performed an analysis of the jitter data collected by the MBA program using our data pipeline to examine jitter by technology between 2011 – 2023. The results for upstream and downstream jitter at the 95th percentile for all measurements, and by technology are shown in Figure III.5. These results indicate all jitter measurements at the 99th percentile or below were less than 10 ms. Jitter thresholds provide guidance that it should be kept below 30 ms to ensure VoIP quality.²⁰ Given these results, it is clear the MBA data reports excellent jitter performance throughout its full history.

¹⁹ To measure jitter, the FCC employs the Packet Delay Variation (PDV) approach described in [RFC5481]. PDV is defined as the difference between the one-way delay of a packet and the minimum one-way delay observed in the transmission path. These tests were executed multiple times daily, and the 99th percentile of results aggregated to assess jitter performance over time. See A. Morton, "Packet Delay Variation Applicability Statement," RFC 5481, IETF, March 2009. Available: <https://doi.org/10.17487/RFC5481>.

²⁰ See Cisco Systems, *Enterprise QoS Solution Reference Network Design Guide*, San Jose, CA: Cisco Systems, 2004. Accessed on June 15, 2025 at <https://community.cisco.com/legacyfs/online/legacy/3/2/6/64623-qossrnd.pdf>.

Figure III.5: Upstream and Downstream Jitter (99th percentile, DSL, Cable, Fiber)



IV. TECHNOLOGY POLICY IMPLICATIONS

Our study finds several notable results from our analysis of the speed, latency, and jitter, measurements obtained over the entire duration of the FCC's MBA program, while highlighting the limitations of interpreting MBA data in our discussion. The following discussion in this section is broken into three different subject matter categories: 1) insights from the longitudinal analysis, 2) insights into the relationship between speed and latency, and 3) broadband policy implications.

Insights from Longitudinal Analysis

The performance data collected for LUL and download/upload speed metrics provide a quantitative measure of the improvement in latency and service speed spanning 2011 – 2023. We recognize that analysis of the MBA Program data cannot support general conclusions regarding overall Internet performance or broadband subscriber experience, but to the extent that the MBA Program provides some representation of trends in Internet performance, imperfect as it may be, it is a useful and encouraging result that there has been substantial improvement in the latency and service speed of residential broadband services over this time based upon performance data obtained from this large sample of major ISPs serving the US.

Starting with the overall trends in download service speed, our analysis calculated the cumulative annual growth rate (CAGR) in the median download speed to be about 20% (1.6 Mbps) for DSL, 30% (22 Mbps) for Cable, and 28% (28 Mbps) for Fiber systems between 2011 - 2023. Notably, between 2020 through July 2023, the CAGR in download speed for Fiber increased to 40% (73 Mbps) for Fiber. For the upstream direction, our analysis found the CAGR of the median upload speed to be about 9% (1.6 Mbps) for DSL, 25% (22 Mbps) for Cable, and 26% (28 Mbps) for Fiber systems between 2011 - 2023. The downstream and upstream CAGR

values between 25% - 30% for Cable and Fiber over the 13-year period of collection demonstrate substantial, sustained growth in downstream and upstream broadband capacity. The incremental nature of the DSL network architecture, which adds upload capacity in smaller increments of capacity and cost in much smaller increments compared to Cable and Fiber, explains the lower upstream CAGR figure.

Regarding latency, our analysis found between 2011 – 2023 that the median downstream LUL decreased an average annual amount of 5 ms (a negative CAGR of 5%) for DSL, 4 ms (-5% CAGR) for Cable, and 10 ms (-13% CAGR) for Fiber. Our analysis found that the median upstream LUL decreased over this same time an average annual amount of 23 ms (-5% CAGR) for DSL, 43 ms for Cable (-17% CAGR), and 7 ms (-14% CAGR) for Fiber. The overall decreasing CAGR trends for Fiber LUL of 13% in the downstream and 14% in the upstream demonstrate the strongest latency performance. Cable has shown sustained performance improvement in the upstream as well with a decrease in CAGR of 17%, though the median value remains higher than Fiber though below DSL. In spite of starting with the highest initial LUL values in downstream and upstream, the CAGR trend of decreasing LUL of 5% in both directions trailed Fiber performance by a wide margin.

The longitudinal IQR plots for all technology types over 2011 – 2023 showed two important characteristics that are worthy of comment:

- All of the plots for downstream and upstream speed demonstrated an IQR that has increased in spread over time between 2011 – 2023 for speed measurements. This implies a rise in the variability of broadband speeds being delivered across the middle 50% of consumers being measured in the panel. In contrast, all of the plots for downstream and upstream LUL demonstrated an IQR that decreased over time for

LuL measurements. This implies improvement during this time in the consistency and reliability of broadband performance during periods of network stress.

- All of the plots for upstream and downstream speed and LUL were skewed right, demonstrating more variability in the 75th percentile values. This implies 75th percentile speed values will be higher (which would be considered a positive for the broadband performance) but the 75th percentile LUL values will be higher as well (which would be considered a negative for broadband performance).

The policy implications for these IVR trends will be discussed further later in this section.

These results have identified differences in the evolving performance across DSL, Cable, and Fiber technologies during the time period of our analysis. Usually, the speed of Fiber was the highest and the latency was lowest, followed by Cable and then DSL. This translates into Fiber networks providing the most reliable support of real-time applications like VoIP and video conferencing. Cable networks are generally acceptable but have a non-negligible fraction of users who may experience degraded experience.

In closing this discussion, we need to caution once again that these results cannot be interpreted to represent a clear characterization of trends in Internet performance as a whole. While our findings are representative of speed and latency performance for thousands of Internet users, we have not weighted the data to reflect actual adoption rates for different speed tiers, which would be a necessary step to better approximate actual Internet usage patterns. In addition, while our analysis was constrained to “off-net” test servers located on the public Internet, we note that many changes in MBA infrastructure were continuous across the multi-year study that inevitably influenced measurement values. Indeed, equipment utilized in any given year varied

between 16 – 182 off-net servers and 4157 – 10,171 Whiteboxes.²¹ Such significant changes to the count of “off-net” servers participating in different years of the study.

Insights Into the Relationship between Speed and Latency

In viewing the results in Section III, we have discussed the inverse relationship between speed and latency and the reasons that increases in transmission speed result in corresponding decreases in the latency of the transmission. The MBA dataset provides a unique opportunity to examine this relationship between speed and latency across the different local access network technologies in most use throughout the US. Such knowledge is useful since policy makers are increasingly establishing speed and latency thresholds for broadband services to ensure high consumer experience but also can have the effect of essentially preclude the use of some technology options for deployment.

Initially, our results focused on speed and latency data only from 2023, and this showed that as broadband download speeds increase from below 50 Mbps to more than 500 Mbps, the median idle latency decreased by over 50% from 35 ms to 15 ms, respectively. When congestion was added to analysis, the downstream and upstream LUL decreased roughly 70% from 90 ms to 25 ms and 85% from 190 to 25 ms, respectively, for this same increase in speed in 2023. From these figures we can estimate the impact of the congestion generated by the MBA speed test was 55 ms above baseline latency for consumers of low-speed tiers below 50 Mbps and 10 ms for consumers of high-speed tiers above 500 Mbps. Generally, these latency levels are unlikely to create noticeable delay in VoIP or web browsing.

Yet we also found instances where the relationship between speed and latency was unpredictable. For example, we did find latency levels that are likely to create significant impacts

²¹ See Schnitzer, Reed, and Gandhi, *op cit*, Table IV.

on the QoE of Internet users in our investigation of the lower percentile results of the downstream and upstream LUL metrics. These results demonstrated that the additional stress put on the network in the form of congestion for the downstream LUL test resulted in latency levels causing noticeable delays for above the 75th percentile for DSL and Cable, and above the 95th percentile for Fiber. The results of the upstream LUL test, however, resulted in latency levels causing noticeable delays for above the 75th percentile for DSL and above the 95th percentile for Fiber as expected, but Cable's performance improved to meeting the 100 ms threshold to the 95th percentile. Given the lower speed and limited capacity of the Cable upstream network, this relative improvement of Cable compared to the other local access technologies was not expected.²²

Our results clearly demonstrate that the higher-speed local access network technologies such as Fiber and Cable exhibit lower latency than lower-speed/higher-latency technologies such as DSL. We know that what favors lower latency with higher speeds is, beyond the lower processing and transmission delay of being able to complete more actions in a unit of time, congestion also decreases due to more payload capacity available for transport. As transmission speeds decrease, latency increases due to longer queuing delays arising due to congestion, slower packet processing in managing data traffic, and additional congestion arising in any shared links.

Yet one must be careful to recognize that latency is not directly proportional to speed in an inverse relationship. That is, our results demonstrate that latency does not scale linearly with speed since high-speed networks can show higher latency due to upstream congestion on a

²² This result may be due to the recent deployment of Advanced Queuing Management (AQM) technologies by Cable that would be reflected by the MBA data in 2023. See F. Baker and G. Fairhurst, "IETF Recommendations Regarding Active Queue Management," RFC 7567, Jul. 2015. Accessed on June 15, 2025, at: <https://www.rfc-editor.org/info/rfc7567>.

shared link or due to excessive queuing in the buffer at the last network interface (known as “bufferbloat”²³). Conversely, a lower-speed fiber plan may still offer low and stable latency.

Our analysis of the MBA dataset confirmed the inverse relationship between speed and latency. We were not able to uncover any predictive mechanisms to this relationship, however, beyond the inverse nature that increases in speed contribute to decreases in latency. This will remain a topic of continuing study to see if any additional insights into the circumstances that permit better forecasting of this relationship beyond what is currently understood. For example, better understanding of the underlying demand for bandwidth and latency and the role of real-time applications requiring better network performance might provide additional useful insights and explanation for the the traffic patterns that are being measured in collection efforts like the MBA program.²⁴

Broadband Policy Implications

We arrange our discussion of broadband policy implications of our preceding analysis into three area of discussion: 1) implications on the definition of broadband, 2) implications of increasing variation in broadband performance, and 3) the value in open data platforms of broadband performance data for data-drive policy.

Implications on the Definition of Broadband

To the extent that the latency measurements collected in the MBA program are representative of real-world trends in residential broadband services, the results of this analysis do have implications for the requirements associated with the definition of broadband services

²³ See P. Harrison, Buffer-bloated router? how to prevent it and improve performance, Commun. ACM, vol. 66, no. 6, p. 73–77, May 2023. Accessed on June 14, 2025, at <https://doi.org/10.1145/3564262>.

²⁴ For example, see Y. Liu, J. Prince, and S. Wallsten, Distinguishing Bandwidth and Latency in Households’ Willingness-to-Pay for Broadband Internet Speed (April 2018). Accessed on June 15, 2025, at: <https://ssrn.com/abstract=2942236>.

associated with government funding and grants for broadband infrastructure. In particular, we applied speed and latency thresholds consistent with current broadband funding requirements in the US (a 100 Mbps/20 Mbps download/upload threshold for speed, and a 100 ms threshold at the 95th percentile for latency) to our analytical results to investigate the circumstances under which the different technologies achieved performance levels establish by these requirements. This analysis indicated the challenges still exist when selecting speed and latency performance requirements to unlock the availability of broadband deployment funds, as these requirements will determine what technologies can be used to build the broadband network.²⁵

DSL technology performance as found by our analysis of the MBA data, for example, was problematic at best in making most of these current thresholds of performance. As DSL technology is viewed credibly today by policy makers and ISPs as being technologically obsolescent, the application of these performance thresholds used in our analysis can be view as appropriate and not designed to favor one technology over another in violation of a technology-neutral policy approach. But it is also true that the adoption of these thresholds has not been supported by a comprehensive broadband performance benchmarking process that accounts for consumer QoE, for example.

Implementation of a transparent methodology to create broadband performance benchmarks that will establish an ongoing Definition of Broadband is an important policy step that has yet to be established by the FCC and the National Telecommunications and Information Administration (NTIA) in the US. This process should establish and rely upon open data

²⁵ While a detailed discussion of the best test measurement methodology and metrics to apply for a definition of broadband is beyond the scope of this analysis, we do note that there is an ongoing discussion of the utility applicability of idle and LUL latency metrics. For the purposes of compliance with latency requirements in the access network, an idle latency test could be viewed as a more transparent and reasonable measure of the actual latency in a broadband network under normal use as compared to the LUL test. See Schnitzer, Reed, and Gandhi, *op cit*, Section V.

collection efforts of broadband performance information like the now concluded FCC MBA program. An open source of broadband performance data is necessary to allow for transparent research and reporting on performance metrics that help to measure the “pulse” of the public Internet.

Implications of Increasing Variation in Broadband Performance

One important finding of our analysis could be the increase in interquartile range (IQR) in the measurement data spanning 2011 to 2023 for broadband speeds. This wide range of values in IQR is due to a wide variance in the delivered speed within each technology panel. Our experience with the MBA dataset suggest that this variance is not within multiple users of a single speed tier but reflects the presence of several different speed tiers provided by each technology type.²⁶ For example, while many associate 1 Gbps broadband with fiber technology, there are in reality other lower speed, and lower priced, speed tiers offered between 50 Mbps and 500 Mbps that are blended together in these figures.

The large magnitude of increase in the IQR that we found has occurred due to the accumulation of more varied service tiers generating a wider spread of speed samples across the set of Whiteboxes over time. More specifically, higher IQR in these broadband performance measures indicates a substantial rise in the variability of broadband speed across the middle 50% of Internet users described by the IQR. The only way to determine if this is an accurate reflection of a similar trend on the greater Internet would be to develop a weighted IQR where the

²⁶ See Table 1 of the 13th MBA Report which lists the different service tiers whose broadband performance was measured in the FCC’s 13th Measuring Broadband America Report. This table includes 19 downstream tiers between 1.5 – 200 Mbps and 16 upstream speed tiers between 513 Kbps to 20 Mbps for DSL, 12 downstream tiers between 50 Mbps to 1 Gbps and 6 upstream speed tiers between 5 – 50 Mbps for Cable, and 10 downstream tiers between 50 Mbps to 1 Gbps and 11 upstream speed tiers between 50 Mbps to 1 Gbps for Fiber. See FCC, 13th Measuring Broadband America Fixed Broadband Report (August 9, 2024). Accessed on June 15, 2025 at <https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-fixed-broadband-thirteenth-report>.

weighting of each user is commensurate with the overall adoption level of the accompanying speed tier. A reasonable hypothesis would be that most consumers are served on the most popular speed tiers that are neither the highest speed (most expensive) offered by an ISP, nor the lowest speed, but instead in the “middle”. If true, this would mean that a weighted IQR would be much narrower than the unweighted IQRs that we have reported in this study.

Nevertheless, the possibility of some increase in the IQRs or speed, even if limited, has several interesting implications for broadband policy that are worth discussing.

First, a steady increase in the download/upload IQR data over time implies greater dispersion or variability in broadband performance as reflected in actual speed measurements. The narrow IQR ranges across DSL, Cable, and Fiber technologies in 2011 implied a relative uniform performance environment across the set of users within each technology type. If by 2023 the IQR range of the middle 50% of users has actually widened for each technology type, then these users are now experiencing a broader range of service speeds that could be significant enough to result in differing QoE.

Such an increase in the heterogeneity in the QoE of consumers by 2023 would actually be capable of generating positive benefits to the market by providing higher speeds to consumers even within the same middle tier or presenting a wider palette of speed choices that allows consumers to select download and upload speeds that are best customized to their personal tastes and desires in broadband service. On the other hand, the heterogeneity could be a reflection of broadening gaps between different regional ISPs where some areas are upgrading broadband infrastructure while others are stagnating. If true, it would raise emerging Digital Divide concerns that more users in the middle range are in the early stages of being left behind relative to those at the top of the broadband performance spectrum. This also would imply technological

and infrastructure gaps are no longer confined to the tail of the distribution in unserved and underserved areas but are becoming embedded in the core requiring eventual attention from broadband funding programs to address.

These are important questions that are raised by our analysis of the MBA dataset that require further research to better understand whether there is additional analysis and data to support this shift or fragmentation of the broadband market. At a minimum, this discussion shows a need for detailed statistical metrics like IQR to assess the extent to which broadband policy is achieving its policy objectives.

Value in Open Data Platforms for Data-Driven Broadband Policy

A critical implication of our analysis is the essential role that open data platforms play in shaping effective, evidence-based broadband policy. The FCC's MBA program demonstrated how transparent, longitudinal data can support public accountability, guide policy design, and inform both industry and academic research. As broadband policy increasingly emphasizes not only access to address the Digital Divide but also performance, reliability, and user experience, open datasets like those generated by the MBA program are foundational to a functional, informed data-driven policy ecosystem.

One of the earliest rationales for the MBA program was to verify whether ISPs were delivering the advertised service speeds. Over time, the program matured into a comprehensive platform for evaluating real-world performance metrics, such as latency under load, jitter, and packet loss—key indicators of broadband quality that are critical to users of real-time applications like video conferencing, telehealth, and cloud gaming. The availability of this open dataset over a 13-year span enabled a unique opportunity to observe long-term trends in both the evolution and variability of broadband service quality.

The longitudinal trends in speed and latency documented in this paper highlight the need for continued collection and public release of similar performance data. Our analysis documents how broadband performance in speed has grown at impressive CAGR rates of 20% - 30% over Cable and Fiber networks between 2011 to 2023. We also documented how latency has declined at CAGR rates between 5% to 13% for Cable and Fiber networks benefitting from the strong increase in speed. The identification of potential increases in IQR by our study raises important policy questions. This nuanced performance landscape cannot be fully captured by simplistic averages or broadband coverage maps. Without open access to standardized, validated datasets like those produced by the MBA, public officials may be left relying on incomplete industry reports when allocating public funds or defining service obligations.

The value of open data is further underscored in an era where broadband funding programs such as the Broadband Equity, Access, and Deployment (BEAD) program are increasingly conditioned on measurable performance benchmarks. These benchmarks must be grounded in a transparent and reproducible methodology for evaluating speed, latency, and other QoE indicators.

Moreover, the open-source data processing pipeline that we developed for this study demonstrates that academic researchers can extend the analytical power of these datasets well beyond their original regulatory purpose. By enriching the MBA data with additional contextual fields, we enabled more granular, policy-relevant analysis of trends in broadband technology performance, and the relationship between speed and latency. These insights would not be possible without public access to the underlying performance data.

Going forward, the sunset of the MBA program in July 2023 leaves a critical gap in the federal broadband data infrastructure. We recommend that policymakers at the FCC and NTIA

prioritize the creation of a successor program that preserves the methodological rigor and public accessibility of the MBA program while expanding its scope to include additional technologies (e.g., fixed wireless, satellite) and emerging metrics (e.g., consistency of service, or real-time application benchmarks such as video quality metrics). This can be accomplished efficiently without having establish significant investment in test clients or Whiteboxes since all ISPs today operate modern network management systems that can capture standard speed, latency, packet loss, and reliability data in sufficient detail for the purposes of any new program.

Without moving forward to replace the MBA program with a successor platform, there is a real risk that broadband policy will be shaped by fragmented, unverified, or selectively disclosed information. A commitment to open, transparent, and independently analyzable performance data is not just a research enabler, it will serve as a cornerstone for pursuing enforcement of broadband policies and accountability for private-sector operations of critical infrastructure to deliver broadband services.